1. A digital data processing system with improved access to information stored on a peripheral device, comprising

A. first and second nodes coupled to one another over a first communications pathway,

B. the second node being coupled to a peripheral device over a second communications pathway,

C. the first node being coupled to the peripheral device over a third communications pathway,

D. a file system, executing on the first and second nodes, being capable of responding to access requests generated by the first node, for

(i) transferring data designated by the request between the first node and the peripheral device via the second node and via the first and second communications pathways,

(ii) maintaining administrative information pertaining to storage of the data designated by the request on the peripheral device,

E. a first bypass, executing on at least the first node, for interceding in response to at least a first selected access request applied thereby to the file system, by transferring data designated by that request between the first node and the peripheral device over the third communications pathway in accord with administrative information maintained by the file system pertaining to storage of that data on the peripheral storage device.

2. A digital data processing system according to claim 1, wherein the transfer of data by the first bypass between the first node and the peripheral device obviates transfer of such data by the file system between the first node and the peripheral device.

3. A digital data processing system according to claim 1, wherein the first bypass transfers the designated data over the third communications pathway to physical storage locations on the peripheral device determined from the administrative information for that data maintained by the second node.

4. A digital data processing system according to claim 3, wherein the first bypass obtains from the second node administrative information from which can be determined physical storage locations in the peripheral device for the data designated by the first request.

5. A digital data processing system according to claim 4, wherein the first bypass at least initiates obtaining administrative information from the second node by generating and applying to the file system a second access request, the second request being for access to a logical unit to which access is controlled by the second node.

6. A digital data processor according to claim 5, wherein the file system transfers the second access request to the second node over the second communications pathway.

7. A digital data processor according to claim 5, wherein the second request is for access to a logical unit other than a file designated in the first access request.

8. A digital data processing system according to claim 7, wherein the logical unit is a file that resides on a peripheral device local to the first node.

9. A digital data processor according to claim 8, wherein the request generated by the first bypass is a request to write and a file.

10. A digital data processing system according to claim 7, comprising a second bypass, executing on the second node, for responding to the second access request by obtaining from the second node, and for at least temporarily retaining, a physical mapping for a file designated by the first access request.

11. A digital data processing system according to claim 10, wherein the first bypass obtains the physical mapping for the file designated by the first access request by generating and applying to the file system a third access request for access, the third request designating the logical unit designated by the second access request.

12. A digital data processing system according to claim 4, wherein

A. the first bypass generates and applies to the file system a second access request, the second access request being a request for writing a logical unit an identifier of a file designated by the first access request, the logical unit being a file to which access is controlled by the second node but which is other than the file designated in the first access request,

B. the second bypass responds to the second access request by reading the logical unit to determine the identifier, obtaining from the second node a physical mapping for a file designated by that identifier, and writing that physical mapping to the logical unit, and

C. the first bypass generates and applies to the file system a third access request for access, the third request being a request for the physical mapping from the logical unit.

13. A digital data processing system according to claim 1, wherein the first bypass selectively limit transfers between the first node and the peripheral device over the third communications pathway.

14. A digital data processing system according to claim 13, wherein the first bypass selectively limits such transfers to facilitate access to the peripheral device by a node other than the first node.

15. A digital data processing system according to claim 14, wherein the bypass limits such transfers in accord with a predefined throttling limit.

16. A digital data processing system according to claim 1, wherein the first communications pathway comprises any of a wide area network interconnect, local area network interconnect, internet interconnect, or other network interconnect.

17 A digital data processing system according to claim 1, wherein the third communications pathway comprises any of a fibre channel, a firewire bus, a serial storage architecture (SSA) bus, high-speed Ethernet bus, high performance parallel interface (HPPI) bus, and other high-speed peripheral device bus.

18. A digital data processing system with improved file access, comprising

A. first and second nodes coupled to one another over a first communications pathway,

B. the second node being coupled to a peripheral device over a second communications pathway,

27

C. the first node being coupled to the peripheral device over a third communications pathway,

D. an applications program executing on the first node generating requests for access to files stored on the peripheral device, and for applying those requests to a file system,

E. the file system, executing on the first and second nodes, being capable of responding to access requests generated by the first node, for

(i) transferring data designated by file access requests between the first node and the peripheral device via the second node and via the first and second communications pathways,

(ii) maintaining administrative information pertaining to physical mappings for files stored on the peripheral device,

E. a first bypass, executing on at least the first node and coupled to the file system, for transferring between the first node and the peripheral device over the third communications pathway data designated by at least selected file access requests, the first bypass transferring such data based on administrative information pertaining to physical mappings maintained by the file system, such transfers being in lieu transfer of that data via the second node and via the first and second communications pathways.

19. A digital data processing system according to claim 18, wherein the first bypass initiates obtaining administrative information pertaining to physical mappings from the second node by generating and applying to the file system a second access request, the

second request being directed to a file for which access is controlled by the second node and which is not a file designated in the first access request.

20. A digital data processing system according to claim 19, wherein the second access request includes an identity of the file designated in the first access request.

21. A digital data processing system according to claim 19, comprising a second bypass, executing on the second node, for responding to the second access request by obtaining from the second node, and for at least temporarily retaining, a physical mapping for a file designated by the first access request.

22. A digital data processing system according to claim 21, wherein the first bypass obtains the physical mapping for the file designated by the first access request by generating and applying to the file system a third access request for access, the third request designating the logical unit designated by the second access request.

23. A digital data processing system according to claim 22, wherein the third access request is directed to the same file as the second access request.

24. A digital data processing system according to claim 19, wherein

A. the first bypass generates and applies to the file system a second access request, the second access request being a request for writing a logical unit an identifier of a file designated by the first access request, the logical unit being a file to which access is controlled by the second node but which is other than the file designated in the first access request,

B. the second bypass responds to the second access request by reading the logical unit to determine the identifier, obtaining from the second node a physical mapping for a file designated by that identifier, and writing that physical mapping to the logical unit, and

C. the first bypass generates and applies to the file system a third access request for access, the third request being a request for the physical mapping from the logical unit.

25. A digital data processing system according to claim 24, wherein at least one of the first bypass stores in a data structure physical mappings for the file designated by the first access request.

26. A digital data processing system according to claim 24, wherein the second bypass flushes any write cache for the file designated by the first access request in connection with obtaining physical mappings for that file from the second node.

27. A scaleable networked digital data processing system with improved access to information stored on a first peripheral device, comprising

A. a plurality of networked nodes, including at least first and second server nodes, coupled to one another via a first communications pathway,

B. the second server node being coupled to a first peripheral device over a second communications pathway,

C. the first server node being coupled to the first peripheral device over a third communications pathway,

30

D. each of the first and second server nodes being coupled to zero, one or more client nodes,

E. a file system, executing on at least the first and second server nodes, being capable of responding to access requests generated by the first server node, for

(i) transferring data designated by the request between the first server node and the first peripheral device via the second server node and via the first and second communications pathways,

(ii) maintaining administrative information pertaining to storage of the data designated by the request on the first peripheral device,

E. a first bypass, executing on at least the first server node, for interceding in response to at least a first selected access request applied thereby to the file system, by transferring data designated by that request between the first server node and the first peripheral device over the third communications pathway in accord with administrative information maintained by the file system pertaining to storage of that data on the peripheral storage device.

28. A ~~scaleable~~ scalable networked digital data processing system according to claim 27, further comprising

A. a third server node coupled to at least the second server node via a fourth communications pathway,

B. the file system, executing on at least the second and third server nodes, being capable of responding to access requests generated by the third server node, for

(i) transferring data designated by the request between the third server node and the first peripheral device via the second server node and via the second and fourth communications pathways,

(ii) maintaining administrative information pertaining to storage of the data designated by the request on the first peripheral device,

C. a second bypass, executing on at least the third server node, for interceding in response to at least selected access requests applied thereby to the file system, by transferring data designated by that request between the third server node and the first peripheral device over the fifth communications pathway in accord with administrative information maintained by the file system pertaining to storage of that data on the peripheral storage device.

29. A method for improved access to a peripheral device in a digital data processing system of the type having

first and second nodes coupled to one another over a first communications pathway,

the second node being coupled to a peripheral device over a second communications pathway,

a file system, executing on the first and second nodes, being capable of responding to access requests generated by the first node, for (i) transferring data designated by the request between the first node and the peripheral device via the second node and via the first and second communications pathways, (ii) maintaining administrative information pertaining to storage of the data designated by the request on the peripheral device,

the method comprising:

A. coupling the first node to the peripheral device over a third communications pathway, and

B. bypassing the file system, in part, in responding to at least a first selected access request applied by the first node to the file system, such bypassing including transferring data designated by that request between the first node and the peripheral device over the third communications pathway in accord with administrative information maintained by the file system pertaining to storage of that data on the peripheral storage device.

30. A method according to claim 29, wherein the bypassing step obviates transfer via the second node and via the first and second communications pathways of data transferred between the first node and the peripheral device over the third communications pathway.

32. A method according to claim 29, wherein the bypassing step comprises transferring the designated data over the third communications pathway to physical storage locations on the peripheral device determined from the administrative information for that data maintained by the second node.

33. A method according to claim 31, wherein the bypassing step comprises at least initiating the obtaining of administrative information from the second node by generating and applying to the file system a second access request, the second request being for access to a logical unit to which access is controlled by the second node.

33
34. A method according to claim 32, wherein the second request is for access to a logical unit other than a file designated in the first access request.

34 33
33. A method according to claim 34, wherein the logical unit is a file that resides on a peripheral device local to the first node.

35 34
36. A method according to claim 35, wherein the request generated by the first bypass is a request to write a file.

Sub A5

36
37. A method according to claim 32, comprising responding to the second access request by obtaining from the second node, and at least temporarily retaining on any of the second node and the peripheral device, a physical mapping for a file designated by the first access request.

37
38. A method according to claim 36, wherein the bypassing step comprises transferring, to the first node, the physical mapping for the file designated by the first access request by generating and applying to the file system a third access request for access, the third request designating the logical unit designated by the second access request.

39
39. A method for improved access to a peripheral device in a scaleable networked digital data processing system of the type having

a plurality of networked nodes, including at least first and second server nodes, coupled to one another via a first communications pathway,

the second server node being coupled to a first peripheral device over a second communications pathway,

each of the first and second server nodes being coupled to zero, one or more client nodes,

a file system, executing on at least the first and second server nodes, being capable of responding to access requests generated by the first server node, for (i) transferring data designated by the request between the first server node and the first peripheral device via the second server node and via the first and second communications pathways, (ii) maintaining administrative information pertaining to storage of the data designated by the request on the first peripheral device,

the method comprising the steps of

A. coupling the first server node to the first peripheral device over a third communications pathway,

B. bypassing the file system, in part, in responding to at least a first selected access request applied by the first server node to the file system, such bypassing including transferring data designated by that request between the first server node and the peripheral device over the third communications pathway in accord with administrative information maintained by the file system pertaining to storage of that data on the peripheral storage device.

40. A method according to claim 39, further wherein the networked digital data processing is of the type having

a third server node coupled to at least the second server node via a fourth communications pathway,

a file system, executing on at least the second and third server nodes, being capable of responding to access requests generated by the third server node, for (i) transferring data designated by the request between the third server node and the first peripheral device via the second server node and via the second and fourth communications pathways, (ii) maintaining administrative information pertaining to storage of the data designated by the request on the first peripheral device,

Sub A5

the method comprising

A. coupling the third server node to the first peripheral device over a fifth communications pathway,

B. bypassing the file system, in part, in responding to at least a first selected access request applied by the third server node to the file system, such bypassing including transferring data designated by that request between the third server node and the peripheral device over the fifth communications pathway in accord with administrative information maintained by the file system pertaining to storage of that data on the peripheral storage device.